arXiv:1402.5056v3 [math.NA] 26 Feb 2014

Efficient arithmetic operations for rank-structured matrices based on hierarchical low-rank updates

Steffen Börm and Knut Reimer

February 27, 2014

Many matrices appearing in numerical methods for partial differential equations and integral equations are *rank-structured*, i.e., they contain submatrices that can be approximated by matrices of low rank. A relatively general class of rank-structured matrices are $\mathcal{H}^2$-matrices: they can reach the optimal order of complexity, but are still general enough for a large number of practical applications.

We consider algorithms for performing algebraic operations with $\mathcal{H}^2$-matrices, i.e., for approximating the matrix product, inverse or factorizations in almost linear complexity. The new approach is based on local low-rank updates that can be performed in linear complexity. These updates can be combined with a recursive procedure to approximate the product of two $\mathcal{H}^2$-matrices, and these products can be used to approximate the matrix inverse and the LR or Cholesky factorization.

Numerical experiments indicate that the new method leads to preconditioners that require $\mathcal{O}(n)$ units of storage, can be evaluated in $\mathcal{O}(n)$ operations, and take $\mathcal{O}(n \log n)$ operations to set up.

Part of this research was funded by the Deutsche Forschungsgemeinschaft in the context of project BO 3289/4-1.

1 Introduction

We consider an elliptic partial differential equation of the form

$$\begin{aligned} -\operatorname{div}\sigma(x)\operatorname{grad}u(x) &= f(x) && \text{for all } x \in \Omega, \\ u(x) &= 0 && \text{for all } x \in \partial\Omega, \end{aligned}$$

where $\Omega \subseteq \mathbb{R}^d$ is a domain and $\sigma : \Omega \to \mathbb{R}^{d\times d}$ assigns each point $x \in \Omega$ a symmetric uniformly positive definite matrix $\sigma(x)$.

Using a Galerkin discretization with a finite element basis $(\varphi_i)_{i\in\mathcal{I}}$ leads to a linear system

$$Ax = b,$$

where the right-hand side $b \in \mathbb{R}^{\mathcal{I}}$ corresponds to f and the solution $x \in \mathbb{R}^{\mathcal{I}}$ represents the approximation of u.

In order to reach a sufficiently accurate approximation, it is usually necessary to work with a large number $n = \#\mathcal{I}$ of basis functions, therefore the linear system can become very large.

We are interested in constructing a good preconditioner $B \in \mathbb{R}^{\mathcal{I}\times\mathcal{I}}$, i.e., a matrix that can be evaluated efficiently and that ensures that iterative solvers like the preconditioned cg iteration [**?**] converge rapidly.

In this paper, we consider preconditioners based on rank-structured matrices. It has been proven [3, 6, 11] that suitably chosen submatrices $A^{-1}|_{\hat{t}\times\hat{s}}$, $\hat{t},\hat{s} \subseteq \mathcal{I}$, of the inverse of A can be approximated by low-rank matrices. The same holds for certain submatrices of the triangular factors L and R of the LR factorization $A = LR$ [2, 14, 11]. These approximations have the remarkable property that they do not depend on the smoothness of the coefficient function σ, even discontinuous coefficients are permitted. Similar results hold for the inverses of matrices resulting from the Galerkin discretization of boundary integral operators [10].

Hierarchical matrices [15, 17, 13, 16] take advantage of this property to construct approximations of A^{-1} in rank-structured representation: a hierarchical decomposition of the index set $\mathcal{I}\times\mathcal{I}$ into appropriate subsets $\hat{t}\times\hat{s}$ is used to determine submatrices that can be approximated by low-rank matrices, and these matrices are represented in factorized form. The result requires only $\mathcal{O}(nk\log n)$ units of storage, where k denotes the rank of the submatrices and depends on the required accuracy.

In order to obtain an approximation of A^{-1}, hierarchical matrices use block-wise constructions based on the multiplication of submatrices and replace the exact matrix products by approximations that can be computed efficiently.

A closer look at the theoretical results [5] suggests that the efficiency of the approximation can be improved by representing the low-rank submatrices using hierarchically nested bases. This leads to $\mathcal{H}^2$-*matrices*, originally developed for integral operators [18, 8]. Using an $\mathcal{H}^2$-matrix representation reduces the storage requirements to $\mathcal{O}(nk)$.

Constructing a good $\mathcal{H}^2$-matrix preconditioner is a challenging task. So far, two approaches have been published: if the nested bases are given a priori, the matrix multiplication can be carried out in $\mathcal{O}(nk^2)$ operations [4], and the inversion and the LR factorization can be reduced to a sequence of matrix multiplications leading to a similar complexity estimate.

In practice, the bases are typically not known in advance, so we have to construct them during the course of our algorithm. The method presented in [7] is able to construct an approximation of the matrix product in $\mathcal{O}(nk^2\log n)$ operations, but requires $\mathcal{O}(nk^2\log^2 n)$ operations for the inverse or the LR factorization, i.e., it is not faster than comparable algorithms for simple hierarchical matrices.

In very special cases, e.g., for essentially one-dimensional problems, HSS-matrices

[9, 21, 24], a special case of $\mathcal{H}^2$-matrices using a very simple block partition, allow us to construct bases on the fly in optimal complexity. Applications of HSS-matrices to two-dimensional geometries rely on special properties to reduce to one-dimensional subproblems [20, 23], while no HSS-algorithms of quasi-linear complexity for three-dimensional problems are currently known.

In this paper, we present a new algorithm for approximating the product of two $\mathcal{H}^2$-matrices. Although this algorithm also requires $\mathcal{O}(nk^2 \log n)$ operations, it can be used to compute LR and Cholesky factorizations as well as an approximate inverse in $\mathcal{O}(nk^2 \log n)$ operations, making it significantly more efficient than previous methods.

The key feature of the new algorithm is the use of efficient local low-rank updates: given an $\mathcal{H}^2$-matrix G, we can add a low-rank matrix $R \in \mathbb{R}^{\hat{t}\times\hat{s}}$ to a submatrix $G|_{\hat{t}\times\hat{s}}$ and ensure that the result is again *globally* an $\mathcal{H}^2$-matrix. This update requires only $\mathcal{O}(k^2(\#\hat{t} + \#\hat{s}))$ operations. Since these update operations are at the heart of most important algorithms for hierarchical matrices, the new algorithm allows us to extend them to $\mathcal{H}^2$-matrices, reducing both the algorithmic complexity and the storage requirements.

The paper is organized as follows: the next section introduces the basic definitions of hierarchical and $\mathcal{H}^2$-matrices. It is followed by a section outlining the ideas of formatted algebraic matrix operations that demonstrates that the most important operations can be reduced to a sequence of local low-rank updates. The next section describes how a low-rank update can be applied efficiently to an $\mathcal{H}^2$-matrix, and it is followed by a section considering the modifications required to perform low-rank updates of submatrices efficiently. The last section is devoted to numerical experiments illustrating that the new algorithm takes $\mathcal{O}(n \log n)$ operations to set up efficient preconditioners in $\mathcal{O}(n)$ units of storage.

2 $\mathcal{H}^2$-matrices

Hierarchical matrices, like most representation schemes for rank-structured matrices, are based on a hierarchical subdivision of the index set. This subdivision can be expressed by a tree structure.

Definition 1 (Cluster tree) *Let $\mathcal{I}$ be a finite index set, and let $\mathcal{T}$ be a labeled tree. Denote the label of each node $t \in \mathcal{T}$ by $\hat{t}$.*

$\mathcal{T}$ is called a cluster tree *for $\mathcal{I}$ if the following conditions hold:*

- *its root $r = \operatorname{root}(\mathcal{T})$ is labeled by $\mathcal{I}$, i.e., $\hat{r} = \mathcal{I}$,*
- *for each $t \in \mathcal{T}$ with $\operatorname{sons}(t) \neq \emptyset$, we have $\hat{t} = \bigcup_{t' \in \operatorname{sons}(t)} \hat{t}'$,*
- *for all $t \in \mathcal{T}$ and $t_1, t_2 \in \operatorname{sons}(t)$ with $t_1 \neq t_2$, we have $\hat{t}_1 \cap \hat{t}_2 = \emptyset$.*

A cluster tree for $\mathcal{I}$ is denoted by $\mathcal{T}_\mathcal{I}$, its nodes are called clusters*, and the set of its leaves is denoted by $\mathcal{L}_\mathcal{I} := \{t \in \mathcal{T}_\mathcal{I} \ : \ \operatorname{sons}(t) = \emptyset\}$.*

We note that the first two conditions imply $\hat{t} \subseteq \mathcal{I}$ for all $t \in \mathcal{T}_\mathcal{I}$, and that all three conditions together imply that the labels of the leaves form a disjoint partition $\{\hat{t} \ : \ t \in \mathcal{L}_\mathcal{I}\}$ of the index set $\mathcal{I}$.

Since we are interested in a partition of a matrix, we construct a special cluster tree for product index sets $\mathcal{I} \times \mathcal{J}$.

Definition 2 (Block tree) *Let $\mathcal{T}_\mathcal{I}$ and $\mathcal{T}_\mathcal{J}$ be cluster trees for index sets $\mathcal{I}$ and $\mathcal{J}$, respectively.*

A labeled tree $\mathcal{T}$ is called a block tree *for $\mathcal{T}_\mathcal{I}$ and $\mathcal{T}_\mathcal{J}$ if the following conditions hold:*

- *for each $b \in \mathcal{T}$, there are $t \in \mathcal{T}_\mathcal{I}$ and $s \in \mathcal{T}_\mathcal{J}$ such that $b = (t, s)$ and $\hat{b} = \hat{t} \times \hat{s}$,*
- *the root $r = \operatorname{root}(\mathcal{T})$ is of the form $r = (\operatorname{root}(\mathcal{T}_\mathcal{I}), \operatorname{root}(\mathcal{T}_\mathcal{J}))$,*
- *for each $b = (t, s) \in \mathcal{T}$ with $\operatorname{sons}(b) \neq \emptyset$, we have*

$$\operatorname{sons}(b) = \begin{cases} \{t\} \times \operatorname{sons}(s) & \text{if } \operatorname{sons}(t) = \emptyset, \operatorname{sons}(s) \neq \emptyset, \\ \operatorname{sons}(t) \times \{s\} & \text{if } \operatorname{sons}(t) \neq \emptyset, \operatorname{sons}(s) = \emptyset, \\ \operatorname{sons}(t) \times \operatorname{sons}(s) & \text{otherwise.} \end{cases}$$

A block tree for $\mathcal{T}_\mathcal{I}$ and $\mathcal{T}_\mathcal{J}$ is denoted by $\mathcal{T}_{\mathcal{I}\times\mathcal{J}}$, its nodes are called blocks*, and the set of its leaves is denoted by $\mathcal{L}_{\mathcal{I}\times\mathcal{J}} := \{b \in \mathcal{T}_{\mathcal{I}\times\mathcal{J}} \ : \ \operatorname{sons}(b) = \emptyset\}$.*

For each block $b = (t, s) \in \mathcal{T}_{\mathcal{I}\times\mathcal{J}}$, we call t the row cluster *(or* target cluster*) and s the* column cluster *(or* source cluster*).*

We note that the definitions imply that a block tree $\mathcal{T}_{\mathcal{I}\times\mathcal{J}}$ is a cluster tree for $\mathcal{I} \times \mathcal{J}$ and that therefore the set $\{\hat{t} \times \hat{s} \ : \ b = (t, s) \in \mathcal{L}_{\mathcal{I}\times\mathcal{J}}\}$ is a disjoint partition of this index set. We use this partition to split matrices into submatrices.

In order to determine which of these submatrices can be approximated by low-rank representations, we split the set $\mathcal{L}_{\mathcal{I}\times\mathcal{J}}$ of leaf blocks into a set $\mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}$ of *admissible* blocks and the remainder $\mathcal{L}^-_{\mathcal{I}\times\mathcal{J}} := \mathcal{L}_{\mathcal{I}\times\mathcal{J}} \setminus \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}$ of *inadmissible* blocks.

We represent the admissible blocks $b = (t, s) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}$ in the form

$$G|_{\hat{t}\times\hat{s}} \approx V_t S_b W_s^*,$$

where the columns of V_t and W_s are interpreted as basis vectors for subsets of $\mathbb{R}^{\hat{t}}$ and $\mathbb{R}^{\hat{s}}$ and S_b contains the coefficients corresponding to these basis vectors. We require the basis vectors to be *hierarchically nested*:

Definition 3 (Cluster basis) *Let $k \in \mathbb{N}$, and let $(V_t)_{t\in\mathcal{T}_\mathcal{I}}$ be a family of matrices satisfying $V_t \in \mathbb{R}^{\hat{t}\times k}$ for all $t \in \mathcal{T}_\mathcal{I}$.*

This family is called a (nested) cluster basis *if for each $t \in \mathcal{T}_\mathcal{I}$ there is a matrix $E_t \in \mathbb{R}^{k\times k}$ such that*

$$V_t|_{\hat{t}'\times k} = V_{t'} E_{t'} \qquad \text{for all } t \in \mathcal{T}_\mathcal{I}, \ t' \in \operatorname{sons}(t). \tag{1}$$

The matrices E_t are called transfer matrices*, and k is called the* rank *of the cluster basis.*

Due to (1), we only have to store the matrices V_t for leaf clusters $t \in \mathcal{L}_\mathcal{I}$ and the transfer matrices E_t for all clusters $t \in \mathcal{T}_\mathcal{I}$. In typical situations, this *nested representation* requires only $\mathcal{O}(n_\mathcal{I} k)$ units of storage [8, 7], where $n_\mathcal{I} := \#\mathcal{I}$ denotes the cardinality of the index set $\mathcal{I}$.

Definition 4 ($\mathcal{H}^2$-matrix) *Let $\mathcal{T}_\mathcal{I}$ and $\mathcal{T}_\mathcal{J}$ be cluster trees for index sets $\mathcal{I}$ and $\mathcal{J}$, let $\mathcal{T}_{\mathcal{I}\times\mathcal{J}}$ be a matching block tree, and let $(V_t)_{t\in\mathcal{T}_\mathcal{I}}$ and $(W_s)_{s\in\mathcal{T}_\mathcal{J}}$ be cluster bases.*

A matrix $G \in \mathbb{R}^{\mathcal{I}\times\mathcal{J}}$ is called an $\mathcal{H}^2$-matrix *for $\mathcal{T}_{\mathcal{I}\times\mathcal{J}}$, $(V_t)_{t\in\mathcal{T}_\mathcal{I}}$ and $(W_s)_{s\in\mathcal{T}_\mathcal{J}}$, if for each admissible block $b = (t, s) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}$ there is a matrix $S_b \in \mathbb{R}^{k\times k}$ such that*

$$G|_{\hat{t}\times\hat{s}} = V_t S_b W_s^*. \tag{2}$$

The matrices S_b are called coupling matrices*, the cluster bases $(V_t)_{t\in\mathcal{T}_\mathcal{I}}$ and $(W_s)_{s\in\mathcal{T}_\mathcal{J}}$ are called* row and column cluster bases*.*

We can represent an $\mathcal{H}^2$-matrix by the cluster bases, the coupling matrices S_b for all admissible blocks $b \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}$ and the *nearfield matrices* $G|_{\hat{t}\times\hat{s}}$ for all inadmissible blocks $b = (t, s) \in \mathcal{L}^-_{\mathcal{I}\times\mathcal{J}}$. In typical situations, this *$\mathcal{H}^2$-matrix representation* requires only $\mathcal{O}((n_\mathcal{I} + n_\mathcal{J})k)$ units of storage [8, 7].

If we want to approximate a given matrix $G \in \mathbb{R}^{\mathcal{I}\times\mathcal{J}}$ by an $\mathcal{H}^2$-matrix $\widetilde{G}$ for cluster bases $(V_t)_{t\in\mathcal{T}_\mathcal{I}}$ and $(W_s)_{s\in\mathcal{T}_\mathcal{J}}$, *orthogonal cluster bases* are very useful: if we assume

$$V_t^* V_t = I, \qquad W_s^* W_s = I \qquad \text{for all } t \in \mathcal{T}_\mathcal{I},\ s \in \mathcal{T}_\mathcal{J},$$

the optimal coupling matrices (with respect both to the Frobenius norm and the spectral norm) are given by

$$S_b := V_t^* G|_{\hat{t}\times\hat{s}} W_s \qquad \text{for all } b = (t, s) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}.$$

This property can be used to compute the best approximation of the product of $\mathcal{H}^2$-matrices in $\mathcal{O}(nk^2)$ operations [4] as long as both cluster bases are known in advance.

3 Algebraic operations

We are looking for a preconditioner for a matrix $A \in \mathbb{R}^{\mathcal{I}\times\mathcal{I}}$ corresponding to a Galerkin discretization of an elliptic partial differential equation or an integral equation. According to [19, 2, 14, 11, 10], we can expect to be able to approximate the matrices $L, R \in \mathbb{R}^{\mathcal{I}\times\mathcal{I}}$ of the standard triangular LR factorization by $\mathcal{H}^2$-matrices.

For the sake of brevity, we restrict our discussion to the case of binary cluster trees, i.e., each cluster has either two sons or none.

Since A is a quadratic matrix, we can use the same cluster tree $\mathcal{T}_\mathcal{I}$ for its rows and columns. We denote the corresponding block tree by $\mathcal{T}_{\mathcal{I}\times\mathcal{I}}$.

The construction of the LR factorization is performed by a recursive procedure: let $t \in \mathcal{T}_\mathcal{I}$. We are looking for the LR factorization of $A|_{\hat{t}\times\hat{t}}$.

If $\operatorname{sons}(t) = \emptyset$, we can assume that $A|_{\hat{t}\times\hat{t}}$ is a small matrix, therefore we can construct the LR factorization by standard Gaussian elimination.

If $\operatorname{sons}(t) = \{t_1, t_2\}$, on the other hand, we split the matrices into submatrices

$$\begin{aligned}
A_{11} &:= A|_{\hat{t}_1 \times \hat{t}_1}, & A_{12} &:= A|_{\hat{t}_1 \times \hat{t}_2}, & A_{21} &:= A|_{\hat{t}_2 \times \hat{t}_1}, & A_{22} &:= A|_{\hat{t}_2 \times \hat{t}_2}, \\
L_{11} &:= L|_{\hat{t}_1 \times \hat{t}_1}, & & & L_{21} &:= L|_{\hat{t}_2 \times \hat{t}_1}, & L_{22} &:= L|_{\hat{t}_2 \times \hat{t}_2}, \\
R_{11} &:= R|_{\hat{t}_1 \times \hat{t}_1}, & R_{12} &:= R|_{\hat{t}_1 \times \hat{t}_2}, & & & R_{22} &:= R|_{\hat{t}_2 \times \hat{t}_2}
\end{aligned}$$

and consider the block equation

$$\begin{aligned}
\begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} &= A|_{\hat{t} \times \hat{t}} = L|_{\hat{t} \times \hat{t}} R|_{\hat{t} \times \hat{t}} \\
&= \begin{pmatrix} L_{11} & \\ L_{21} & L_{22} \end{pmatrix} \begin{pmatrix} R_{11} & R_{12} \\ & R_{22} \end{pmatrix} = \begin{pmatrix} L_{11} R_{11} & L_{11} R_{12} \\ L_{21} R_{11} & L_{22} R_{22} + L_{21} R_{12} \end{pmatrix}.
\end{aligned}$$

We can solve $L_{11} R_{11} = A_{11}$ by recursion and obtain L_{11} and R_{11}. Then we can solve the triangular systems $L_{11} R_{12} = A_{12}$ and $L_{21} R_{11} = A_{21}$ by forward substitution to obtain R_{12} and L_{21}. In a last step, we can solve $L_{22} R_{22} = A_{22} - L_{21} R_{12}$ by recursion to obtain L_{22} and R_{22}, completing the algorithm.

The block forward substitution can also be handled by recursion: in order to solve $LX = Y$ to obtain $X \in \mathbb{R}^{\hat{t} \times \hat{s}}$ given $Y \in \mathbb{R}^{\hat{t} \times \hat{s}}$, we let

$$X_1 := X|_{\hat{t}_1 \times \hat{s}}, \qquad X_2 := X|_{\hat{t}_2 \times \hat{s}}, \qquad Y_1 := Y|_{\hat{t}_1 \times \hat{s}}, \qquad Y_2 := Y|_{\hat{t}_2 \times \hat{s}}$$

and consider

$$\begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix} = Y = L|_{\hat{t} \times \hat{t}} X = \begin{pmatrix} L_{11} & \\ L_{21} & L_{22} \end{pmatrix} \begin{pmatrix} X_1 \\ X_2 \end{pmatrix} = \begin{pmatrix} L_{11} X_1 \\ L_{21} X_1 + L_{22} X_2 \end{pmatrix}.$$

We can solve $L_{11} X_1 = Y_1$ by recursion and obtain X_1. In the second step, we solve $L_{22} X_2 = Y_2 - L_{21} X_1$ again by recursion and obtain X_2, completing the algorithm.

For the system $XR = Y$ with $X, Y \in \mathbb{R}^{\hat{s} \times \hat{t}}$, we can follow a similar approach: we let

$$X_1 := X|_{\hat{s} \times \hat{t}_1}, \qquad X_2 := X|_{\hat{s} \times \hat{t}_2}, \qquad Y_1 := Y|_{\hat{s} \times \hat{t}_1}, \qquad Y_2 := Y|_{\hat{s} \times \hat{t}_2}$$

and arrive at the block equation

$$\begin{pmatrix} Y_1 & Y_2 \end{pmatrix} = Y = XR|_{\hat{t} \times \hat{t}} = \begin{pmatrix} X_1 & X_2 \end{pmatrix} \begin{pmatrix} R_{11} & R_{12} \\ & R_{22} \end{pmatrix} = \begin{pmatrix} X_1 R_{11} & X_1 R_{12} + X_2 R_{22} \end{pmatrix},$$

requiring us to solve $X_1 R_{11} = Y_1$ and $X_2 R_{22} = Y_2 - X_1 R_{12}$ by recursion.

We conclude that we only need an efficient algorithm for computing the matrix product in order to derive efficient algorithms for the block forward substitution and the LR factorization. The matrix inverse can be constructed by a similar procedure.

Both the inversion and the LR factorization require updates to submatrices of the form

$$Z|_{\hat{t} \times \hat{r}} \leftarrow Z|_{\hat{t} \times \hat{r}} + \alpha X|_{\hat{t} \times \hat{s}} Y|_{\hat{s} \times \hat{r}}$$

with $(t, r), (t, s), (s, r) \in \mathcal{T}_{\mathcal{I} \times \mathcal{J}}$. If both (t, s) and (s, r) are not leaves, we can use recursion once more. For the sake of brevity, we consider only the most general case that none

of the clusters is a leaf, i.e., that we have $\operatorname{sons}(t) = \{t_1, t_2\}$, $\operatorname{sons}(s) = \{s_1, s_2\}$ and $\operatorname{sons}(r) = \{r_1, r_2\}$. As before, we split $X|_{\hat{t}\times\hat{s}}$, $Y|_{\hat{s}\times\hat{r}}$ and $Z|_{\hat{t}\times\hat{r}}$ into submatrices

$$X|_{\hat{t}\times\hat{s}} = \begin{pmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{pmatrix}, \qquad Y|_{\hat{s}\times\hat{r}} = \begin{pmatrix} Y_{11} & Y_{12} \\ Y_{21} & Y_{22} \end{pmatrix}, \qquad Z|_{\hat{t}\times\hat{r}} = \begin{pmatrix} Z_{11} & Z_{12} \\ Z_{21} & Z_{22} \end{pmatrix}$$

and obtain

$$Z|_{\hat{t}\times\hat{r}} + \alpha X|_{\hat{t}\times\hat{s}} Y|_{\hat{s}\times\hat{r}} = \begin{pmatrix} Z_{11} + \alpha X_{11}Y_{11} + \alpha X_{12}Y_{21} & Z_{12} + \alpha X_{11}Y_{12} + \alpha X_{12}Y_{22} \\ Z_{21} + \alpha X_{21}Y_{11} + \alpha X_{22}Y_{21} & Z_{22} + \alpha X_{21}Y_{12} + \alpha X_{22}Y_{22} \end{pmatrix}.$$

We conclude that the product can be computed by recursively carrying out the updates

$$\begin{aligned} Z_{11} &\leftarrow Z_{11} + \alpha X_{11}Y_{11}, & Z_{12} &\leftarrow Z_{12} + \alpha X_{11}Y_{12}, \\ Z_{21} &\leftarrow Z_{21} + \alpha X_{21}Y_{11}, & Z_{22} &\leftarrow Z_{22} + \alpha X_{21}Y_{12}, \\ Z_{11} &\leftarrow Z_{11} + \alpha X_{12}Y_{21}, & Z_{12} &\leftarrow Z_{12} + \alpha X_{12}Y_{22}, \\ Z_{21} &\leftarrow Z_{21} + \alpha X_{22}Y_{21}, & Z_{22} &\leftarrow Z_{22} + \alpha X_{22}Y_{22}. \end{aligned}$$

If (t, r) is a leaf, we can temporarily create submatrices Z_{11}, Z_{12}, Z_{21} and Z_{22} to receive the result of the recursive procedure. These submatrices can then be extended with zeros and added to the block $Z|_{\hat{t}\times\hat{r}}$.

If (t, s) or (s, r) is a leaf block, the term $X|_{\hat{t}\times\hat{s}} Y|_{\hat{s}\times\hat{r}}$ is a low-rank matrix, and a factorized representation can be obtained easily by multiplying the factorized representation of the low-rank matrix with the other matrix [8, Sections 3.2 and 3.3].

In both cases, we need an efficient algorithm that adds a low-rank matrix in factorized representation to an existing $\mathcal{H}^2$-matrix. For hierarchical matrices, this task can be handled by a relatively simple approach: since all low-rank blocks are independent, we only have to add a low-rank matrix to another low-rank matrix, and a singular value decomposition can be used to reduce the rank of the resulting matrix.

For $\mathcal{H}^2$-matrices, adding a low-rank matrix is considerably more challenging, since the blocks shared row and column cluster bases, so changing one block changes the entire corresponding block row and column.

4 Recompression

We first consider a *global* low-rank update. Let $Z \in \mathbb{R}^{\mathcal{I}\times\mathcal{J}}$ be an $\mathcal{H}^2$-matrix with row cluster basis $(V_t)_{t\in\mathcal{T}_\mathcal{I}}$ and column cluster basis $(W_s)_{s\in\mathcal{T}_\mathcal{J}}$ for a block tree $\mathcal{T}_{\mathcal{I}\times\mathcal{J}}$.

Let $X \in \mathbb{R}^{\mathcal{I}\times k}$ and $Y \in \mathbb{R}^{\mathcal{J}\times k}$. We are looking for an efficient algorithm for approximating the sum $\widetilde{Z} := Z + XY^*$.

Our starting point is a simple observation: for each admissible block $b = (t, s) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}$, we have

$$\begin{aligned} (Z + XY^*)|_{\hat{t}\times\hat{s}} &= Z|_{\hat{t}\times\hat{s}} + X|_{\hat{t}\times k} Y|^*_{\hat{s}\times k} = V_t S_b W_s^* + X|_{\hat{t}\times k} I Y|^*_{\hat{s}\times k} \\ &= \begin{pmatrix} V_t & X|_{\hat{t}\times k} \end{pmatrix} \begin{pmatrix} S_b & \\ & I \end{pmatrix} \begin{pmatrix} W_s & Y|_{\hat{s}\times k} \end{pmatrix}^*, \end{aligned}$$

i.e., all submatrices are already given in factorized form. By introducing new cluster bases and coupling matrices

$$\widetilde{V}_t := \begin{pmatrix} V_t & X|_{\hat{t}\times k} \end{pmatrix}, \qquad \text{for all } t \in \mathcal{T}_{\mathcal{I}}, \tag{3a}$$

$$\widetilde{W}_s := \begin{pmatrix} W_s & Y|_{\hat{s}\times k} \end{pmatrix}, \qquad \text{for all } s \in \mathcal{T}_{\mathcal{J}}, \tag{3b}$$

$$\widetilde{S}_b := \begin{pmatrix} S_b & \\ & I \end{pmatrix} \qquad \text{for all } b \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}, \tag{3c}$$

we have found an *exact* $\mathcal{H}^2$-matrix representation of $\widetilde{Z} = Z + XY^*$. The transfer matrices for $(\widetilde{V}_t)_{t\in\mathcal{T}_{\mathcal{I}}}$ and $(\widetilde{W}_s)_{s\in\mathcal{T}_{\mathcal{J}}}$ can be easily obtained by extending the transfer matrices of the original cluster bases with the identity matrix, similar to the way the coupling matrices S_b are extended to $\widetilde{S}_b$.

Unfortunately the rank of this representation $\widetilde{Z}$ has doubled compared to the original $\mathcal{H}^2$-matrix Z, therefore using it repeatedly as required by the matrix multiplication would lead quickly to very large ranks and a very inefficient algorithm.

We can fix this problem by using a *recompression algorithm* that takes an $\mathcal{H}^2$-matrix and approximates it by an $\mathcal{H}^2$-matrix of lower rank. This closely resembles the truncation operation that is at the heart of algebraic operations for hierarchical matrices.

An appropriate algorithm is described in [7, Section 6.6]: we aim to construct an improved orthogonal row cluster basis $(Q_t)_{t\in\mathcal{T}_{\mathcal{I}}}$. A column cluster basis can be handled similarly, replacing $\widetilde{Z}$ by $\widetilde{Z}^*$. The matrix Q_t should be chosen in such a way that all blocks connected to the cluster t or one of its predecessors

$$\operatorname{pred}(t) := \begin{cases} \{t\} \cup \operatorname{pred}(t^+) & \text{if there is a } t^+ \in \mathcal{T}_{\mathcal{I}} \text{ with } t \in \operatorname{sons}(t^+), \\ \{t\} & \text{otherwise, i.e., if } t = \operatorname{root}(\mathcal{T}_{\mathcal{I}}) \end{cases}$$

can be approximated in the range of Q_t. If we denote the set of all these clusters by

$$\operatorname{row}^*(t) := \bigcup_{t^*\in\operatorname{pred}(t)} \operatorname{row}(t^*), \qquad \operatorname{row}(t) := \{s \in \mathcal{T}_{\mathcal{J}} \ : \ (t,s) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}\},$$

and take advantage of the fact that we are looking for an orthogonal cluster basis, we have to ensure

$$Q_t Q_t^* \widetilde{Z}|_{\hat{t}\times\hat{s}} \approx \widetilde{Z}|_{\hat{t}\times\hat{s}} \qquad \text{for all } s \in \operatorname{row}^*(t).$$

The task of finding an orthogonal basis that approximates multiple matrices simultaneously can be simplified by combining all of these matrices in a large matrix: we let $\tau := \#\operatorname{row}^*(t)$ and $\operatorname{row}^*(t) = \{s_1, \ldots, s_\tau\}$ and introduce the matrix

$$\widetilde{Z}_t := \begin{pmatrix} \widetilde{Z}|_{\hat{t}\times\hat{s}_1} & \ldots & \widetilde{Z}|_{\hat{t}\times\hat{s}_\tau} \end{pmatrix}.$$

If we have found an orthogonal matrix Q_t of low rank satisfying

$$Q_t Q_t^* \widetilde{Z}_t \approx \widetilde{Z}_t, \tag{4}$$

we have solved our problem. Applying the singular value decomposition directly to this task would lead to an algorithm of at least quadratic complexity.

Fortunately, we can take advantage of the fact that $\widetilde{Z}_t$ is an $\mathcal{H}^2$-matrix: for each $s \in \operatorname{row}(t)$, we have $(t,s) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}$ by definition and therefore

$$\widetilde{Z}|_{\hat{t}\times\hat{s}} = \widetilde{V}_t \widetilde{S}_b \widetilde{W}_s^* = \widetilde{V}_t B_{ts}^*$$

with $B_{ts} := \widetilde{W}_s \widetilde{S}_b^* \in \mathbb{R}^{\hat{s}\times(2k)}$. Since the cluster basis $(\widetilde{V}_t)_{t\in\mathcal{T}_\mathcal{I}}$ is nested, we can extend this result: for each $s \in \operatorname{row}^*(t)$, we find a matrix $B_{ts} \in \mathbb{R}^{\hat{s}\times(2k)}$ with

$$\widetilde{Z}|_{\hat{t}\times\hat{s}} = \widetilde{V}_t B_{ts}^*,$$

therefore we have

$$\widetilde{Z}_t = \begin{pmatrix} \widetilde{Z}|_{\hat{t}\times\hat{s}_1} & \dots & \widetilde{Z}|_{\hat{t}\times\hat{s}_\tau} \end{pmatrix} = \begin{pmatrix} \widetilde{V}_t B_{t,s_1}^* & \dots & \widetilde{V}_t B_{t,s_\tau}^* \end{pmatrix} = \widetilde{V}_t B_t^*, \quad B_t := \begin{pmatrix} B_{t,s_1} \\ \vdots \\ B_{t,s_\tau} \end{pmatrix}. \tag{5}$$

We conclude that the matrices $\widetilde{Z}_t$ are of low rank and given in factorized form.

If we solve (4) by computing the singular value decomposition of $\widetilde{Z}_t$, we only require the left singular vectors and the singular values to construct Q_t. This means that applying orthogonal transformations to the columns of $\widetilde{Z}_t$ will not change the result of our algorithm. We can take advantage of this property to reduce the computational work: let $P_t \widetilde{B}_t = B_t$ be a QR decomposition of B_t with $\widetilde{B}_t \in \mathbb{R}^{(2k)\times(2k)}$ and an orthogonal matrix P_t. Replacing B_t^* by $B_t^* P_t$ in (5) yields that (4) is equivalent to

$$Q_t Q_t^* \widetilde{V}_t \widetilde{B}_t^* \approx \widetilde{V}_t \widetilde{B}_t^*, \tag{6}$$

so Q_t can be obtained by computing the singular value decomposition of the matrix $\widetilde{V}_t \widetilde{B}_t^*$ with only $2k$ columns.

Finding the QR decomposition $P_t \widetilde{B}_t = B_t$ directly would again lead to an algorithm of quadratic complexity, but we can once more take advantage of the properties of $\mathcal{H}^2$-matrices: the cluster basis $(\widetilde{V}_t)_{t\in\mathcal{T}_\mathcal{I}}$ is nested. If we have $s \in \operatorname{row}^*(t)$ with $(t,s) \notin \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}$, the definition implies $s \in \operatorname{row}^*(t^+)$, where t^+ denotes the father of t. Therefore we can find $B_{t^+,s} \in \mathbb{R}^{\hat{s}\times(2k)}$ such that

$$\widetilde{Z}|_{\hat{t}^+\times\hat{s}} = \widetilde{V}_{t^+} B_{t^+,s}^*,$$

and (1) yields

$$\widetilde{Z}|_{\hat{t}\times\hat{s}} = \widetilde{V}_{t^+}|_{\hat{t}\times k} B_{t^+,s}^* = \widetilde{V}_t \widetilde{E}_t B_{t^+,s}^*,$$

where $\widetilde{E}_t$ denotes the transfer matrix corresponding to the cluster t.

This means that all blocks connected to "strict" predecessors of t are also present in B_{t^+}, and at least part of the QR decomposition of B_t can be inherited from the father t^+. Let $\varrho := \#\operatorname{row}(t)$ and $\{s_1,\dots,s_\varrho\} := \operatorname{row}(t)$. Then we can write B_t as

$$B_t = \begin{pmatrix} B_{t,s_1} \\ \vdots \\ B_{t,s_\varrho} \\ B_{t^+}\widetilde{E}_t^* \end{pmatrix} = \begin{pmatrix} B_{t,s_1} \\ \vdots \\ B_{t,s_\varrho} \\ P_{t^+}\widetilde{B}_{t^+}\widetilde{E}_t^* \end{pmatrix} = \begin{pmatrix} I & & & \\ & \ddots & & \\ & & I & \\ & & & P_{t^+} \end{pmatrix} \begin{pmatrix} B_{t,s_1} \\ \vdots \\ B_{t,s_\varrho} \\ \widetilde{B}_{t^+}\widetilde{E}_t^* \end{pmatrix}.$$

The left factor is already orthogonal, therefore finding the QR decomposition of B_t only requires us to find the decomposition of the right factor. This recursive approach significantly reduces the number of rows, since $\widetilde{B}_{t^+}\widetilde{E}_t^*$ is only a small $(2k) \times (2k)$ matrix, and leads to an algorithm that computes all weight matrices $(\widetilde{B}_t)_{t\in\mathcal{T}_\mathcal{I}}$ in $\mathcal{O}(n_\mathcal{I}k^2)$ operations [7, Algorithm 28].

For an $\mathcal{H}^2$-matrix representation, we require a *nested* cluster basis, and we briefly summarize an algorithm [7, Algorithm 30] that directly computes the transfer matrices: if t is a leaf, we compute the singular value decomposition of $\widetilde{V}_t\widetilde{B}_t^*$ directly and use the left singular vectors corresponding to the largest singular values to construct Q_t.

If t is not a leaf, we have $\operatorname{sons}(t) = \{t_1, t_2\}$ and compute Q_{t_1} and Q_{t_2} by recursion. Since we are looking for a nested cluster basis, we only have to construct transfer matrices F_{t_1} and F_{t_2} with

$$Q_t = \begin{pmatrix} Q_{t_1}F_{t_1} \\ Q_{t_2}F_{t_2} \end{pmatrix} = \begin{pmatrix} Q_{t_1} & \\ & Q_{t_2} \end{pmatrix} \widehat{Q}_t, \qquad \widehat{Q}_t := \begin{pmatrix} F_{t_1} \\ F_{t_2} \end{pmatrix}. \tag{7}$$

This leads to

$$\begin{pmatrix} Q_{t_1} & \\ & Q_{t_2} \end{pmatrix} \widehat{Q}_t\widehat{Q}_t^* \begin{pmatrix} Q_{t_1}^* & \\ & Q_{t_2}^* \end{pmatrix} \widetilde{V}_t\widetilde{B}_t^* = Q_tQ_t^*\widetilde{V}_t\widetilde{B}_t^* \approx \widetilde{V}_t\widetilde{B}_t^*,$$

and multiplying both sides with the adjoints of $Q_{t_1}^*$ and $Q_{t_2}^*$ yields

$$\widehat{Q}_t\widehat{Q}_t^* \begin{pmatrix} Q_{t_1}^*\widetilde{V}_t|_{\hat{t}_1\times k} \\ Q_{t_2}^*\widetilde{V}_t|_{\hat{t}_2\times k} \end{pmatrix} \widetilde{B}_t^* \approx \begin{pmatrix} Q_{t_1}^* & \\ & Q_{t_2}^* \end{pmatrix} \widetilde{V}_t\widetilde{B}_t^* = \begin{pmatrix} Q_{t_1}^*\widetilde{V}_t|_{\hat{t}_1\times k} \\ Q_{t_2}^*\widetilde{V}_t|_{\hat{t}_2\times k} \end{pmatrix} \widetilde{B}_t^*.$$

Since $(\widetilde{V}_t)_{t\in\mathcal{T}_\mathcal{I}}$ is a *nested* cluster basis, we can use (1) to obtain

$$\widehat{V}_t := \begin{pmatrix} Q_{t_1}^*\widetilde{V}_t|_{\hat{t}_1\times k} \\ Q_{t_2}^*\widetilde{V}_t|_{\hat{t}_2\times k} \end{pmatrix} = \begin{pmatrix} Q_{t_1}^*\widetilde{V}_{t_1}\widetilde{E}_{t_1} \\ Q_{t_2}^*\widetilde{V}_{t_2}\widetilde{E}_{t_2} \end{pmatrix}$$

and arrive at the approximation

$$\widehat{Q}_t\widehat{Q}_t^*\widehat{V}_t\widetilde{B}_t^* \approx \widehat{V}_t\widetilde{B}_t^*.$$

It is of the same form as (6), and we can again use the singular value decomposition to construct the matrix $\widehat{Q}_t$. Splitting $\widehat{Q}_t$ according to (7) yields the required transfer matrices. Since $\widehat{V}_t\widetilde{B}_t^*$ is only a $(2k) \times (2k)$ matrix, the recursive algorithm requires only $\mathcal{O}(nk^2)$ operations [7, Theorem 6.27] to construct the entire cluster basis $(Q_t)_{t\in\mathcal{T}_\mathcal{I}}$.

This algorithm also provides us with matrices $R_t = Q_t^*\widetilde{V}_t = \widehat{Q}_t^*\widehat{V}_t$ that can be used to compute the coupling matrices for the recompressed $\mathcal{H}^2$-matrix efficiently: since Q_t is an orthogonal matrix, the best approximation of $V_tS_bW_s^*$ in the new basis is given by the orthogonal projection $Q_tQ_t^*V_tS_bW_s^* = Q_tR_tS_bW_s^*$, so a multiplication of S_b with the small matrix R_t is sufficient to switch a block to the new basis. Applying this procedure to the entire matrix takes $\mathcal{O}(nk^2)$ operations.

By including suitable weighting factors in the matrices B_{ts}, we can guarantee that the recompressed matrix satisfies relative or absolute error tolerances in the Frobenius or spectral norm [7, Section 6.8].

To summarize: the matrix $\widetilde{Z} = Z + XY^*$ can be expressed as an $\mathcal{H}^2$-matrix of rank $2k$ using (3), and the recompression algorithm can be used to construct an approximation of reduced rank in $\mathcal{O}((n_{\mathcal{I}} + n_{\mathcal{J}})k^2)$ operations. This means that *global* low-rank updates can be carried out in linear complexity.

5 Local updates

The multiplication algorithm applies low-rank updates to *submatrices*, not to the entire matrix, therefore we have to modify the algorithm outlined in the previous section.

Let us assume that we want to add a low-rank matrix XY^* with $X \in \mathbb{R}^{\hat{t}_0 \times k}$ and $Y \in \mathbb{R}^{\hat{s}_0 \times k}$ to a matrix block $Z|_{\hat{t}_0 \times \hat{s}_0}$ with $b_0 = (t_0, s_0) \in \mathcal{T}_{\mathcal{I}\times\mathcal{J}}$ and leave the remainder of the matrix Z essentially unchanged. A simple algorithm would be to extend X and Y by zero and use the algorithm presented before, but this would lead to a relatively high computational complexity.

If Z is an $\mathcal{H}^2$-matrix, the submatrix $Z|_{\hat{t}_0 \times \hat{s}_0}$ is also an $\mathcal{H}^2$-matrix with cluster bases and coupling matrices taken from subtrees: let $\mathcal{T}_{\hat{t}_0}$ denote the subtree of $\mathcal{T}_{\mathcal{I}}$ with root t_0, $\mathcal{T}_{\hat{s}_0}$ the subtree of $\mathcal{T}_{\mathcal{J}}$ with root s_0, and $\mathcal{T}_{\hat{t}_0 \times \hat{s}_0}$ the subtree of $\mathcal{T}_{\mathcal{I}\times\mathcal{J}}$ with root $b_0 = (t_0, s_0)$. $Z|_{\hat{t}_0 \times \hat{s}_0}$ is an $\mathcal{H}^2$-matrix for the block tree $\mathcal{T}_{\hat{t}_0 \times \hat{s}_0}$ with row basis $(V_t)_{t \in \mathcal{T}_{\hat{t}_0}}$, column basis $(W_s)_{s \in \mathcal{T}_{\hat{s}_0}}$. Therefore we can apply the algorithm given above and update the block in $\mathcal{O}((\#\hat{t}_0 + \#\hat{s}_0)k^2)$ operations.

Unfortunately, the result will in general no longer be an $\mathcal{H}^2$-matrix, since $Z|_{(\mathcal{I}\setminus\hat{t}_0)\times\hat{s}_0}$ and $Z|_{\hat{t}_0\times(\mathcal{J}\setminus\hat{s}_0)}$ would still be represented by the original cluster bases, not by the ones constructed for the update.

We can fix this problem by switching *all* blocks to the new cluster bases: Applying the update to $Z|_{\hat{t}_0 \times \hat{s}_0}$ yields new cluster bases $(Q_t)_{t \in \mathcal{T}_{\hat{t}_0}}$ and $(Q_s)_{s \in \mathcal{T}_{\hat{s}_0}}$ and, as mentioned above, also matrices $R_t = Q_t^* \widetilde{V}_t$ and $R_s = Q_s^* \widetilde{W}_s$ describing the change of basis. Using these matrices, we can update the coupling matrices:

$$
\begin{aligned}
S_b &\leftarrow R_t \begin{pmatrix} S_b & 0 \\ 0 & I \end{pmatrix} R_s^* && \text{for all } b = (t,s) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}},\ t \in \mathcal{T}_{\hat{t}_0},\ s \in \mathcal{T}_{\hat{s}_0},\\
S_b &\leftarrow R_t \begin{pmatrix} S_b \\ 0 \end{pmatrix} && \text{for all } b = (t,s) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}},\ t \in \mathcal{T}_{\hat{t}_0},\ s \notin \mathcal{T}_{\hat{s}_0},\\
S_b &\leftarrow \begin{pmatrix} S_b & 0 \end{pmatrix} R_s^* && \text{for all } b = (t,s) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}},\ t \notin \mathcal{T}_{\hat{t}_0},\ s \in \mathcal{T}_{\hat{s}_0}.
\end{aligned}
$$

In order to obtain an $\mathcal{H}^2$-matrix, we also have to update the cluster bases $(V_t)_{t \in \mathcal{T}_{\mathcal{I}}}$ and $(W_s)_{s \in \mathcal{T}_{\mathcal{J}}}$. Thanks to the nested representation of these bases, this update is particularly simple: we replace the original cluster bases in the relevant subtrees by the new bases

$$
\begin{aligned}
V_t &\leftarrow Q_t && \text{for all } t \in \mathcal{T}_{\hat{t}_0},\\
W_s &\leftarrow Q_s && \text{for all } s \in \mathcal{T}_{\hat{s}_0}.
\end{aligned}
$$

and update the transfer matrices by

$$E_{t_0} \leftarrow R_{t_0} \begin{pmatrix} E_{t_0} \\ 0 \end{pmatrix}, \qquad E_{s_0} \leftarrow R_{s_0} \begin{pmatrix} E_{s_0} \\ 0 \end{pmatrix}$$

in order to ensure that the resulting cluster bases are still properly nested.

We do not have to change coupling matrices connected to proper predecessors of t_0 or s_0, since the corresponding cluster bases implicitly inherit the update via the transfer matrices E_{t_0} and E_{s_0}.

In order to ensure that the new cluster bases are able to approximate blocks outside of $\mathcal{T}_{\hat{t}_0 \times \hat{s}_0}$, we have to make sure that the weight matrices $\widetilde{B}_t$ are computed correctly: they have to take *all* blocks $(t, s) \in \mathcal{L}^+_{\mathcal{I}\times\mathcal{J}}$ with $s \in \operatorname{row}(t)$ into account, not only the blocks in $\mathcal{T}_{\hat{t}_0 \times \hat{s}_0}$.

Applying [7, Lemma 6.26] to the subtrees $\mathcal{T}_{\hat{t}_0}$ and $\mathcal{T}_{\hat{s}_0}$ yields that the construction of the weight matrices takes only $\mathcal{O}((\#\hat{t}_0 + \#\hat{s}_0)k^2)$ operations. Updating the global cluster bases means only changing the two transfer matrices E_{t_0} and E_{s_0}, and this can be accomplished in $\mathcal{O}(k^3)$ operations. Updating the coupling matrices requires not more than

$$\sum_{\substack{(t,s)\in\mathcal{L}^+_{\mathcal{I}\times\mathcal{J}} \\ t\in\mathcal{T}_{\hat{t}_0}}} (2k)^3 + \sum_{\substack{(t,s)\in\mathcal{L}^+_{\mathcal{I}\times\mathcal{J}} \\ s\in\mathcal{T}_{\hat{s}_0}}} (2k)^3 \leq 8C_{\mathrm{sp}} \sum_{t\in\mathcal{T}_{\hat{t}_0}} k^3 + 8C_{\mathrm{sp}} \sum_{s\in\mathcal{T}_{\hat{s}_0}} k^3 = 8C_{\mathrm{sp}}(\#\mathcal{T}_{\hat{t}_0} + \#\mathcal{T}_{\hat{s}_0})k^3$$

operations, where C_{sp} is the sparsity constant [8] of the block tree $\mathcal{T}_{\mathcal{I}\times\mathcal{J}}$ satisfying

$$\begin{aligned} \#\{s \in \mathcal{T}_{\mathcal{J}} \ : \ (t,s) \in \mathcal{T}_{\mathcal{I}\times\mathcal{J}}\} &\leq C_{\mathrm{sp}} && \text{for all } t \in \mathcal{T}_{\mathcal{I}}, \\ \#\{t \in \mathcal{T}_{\mathcal{I}} \ : \ (t,s) \in \mathcal{T}_{\mathcal{I}\times\mathcal{J}}\} &\leq C_{\mathrm{sp}} && \text{for all } s \in \mathcal{T}_{\mathcal{J}}. \end{aligned}$$

With the standard assumptions $k^3\#\mathcal{T}_{\hat{t}_0} \in \mathcal{O}(k^2\#\hat{t}_0)$ and $k^3\#\mathcal{T}_{\hat{s}_0} \in \mathcal{O}(k^2\#\hat{s}_0)$, we conclude that these updates also take no more than $\mathcal{O}((\#\hat{t}_0 + \#\hat{s}_0)k^2)$ operations.

In summary, we have proven the following complexity estimate:

Theorem 5 (Complexity) *Approximating the low-rank update $Z|_{\hat{t}_0\times\hat{s}_0} \leftarrow Z|_{\hat{t}_0\times\hat{s}_0} + XY^*$ of a submatrix with $b_0 = (t_0, s_0) \in \mathcal{T}_{\mathcal{I}\times\mathcal{J}}$ requires $\mathcal{O}((\#\hat{t}_0 + \#\hat{s}_0)k^2)$ operations.*

6 Complexity of the LR factorization

Having established an estimate for the complexity of low-rank updates, we can now turn our attention towards complexity estimates for the matrix multiplication, forward substitution and the LR factorization.

Let $t, s, r \in \mathcal{T}_{\mathcal{I}}$ such that $(t, s) \in \mathcal{T}_{\mathcal{I}\times\mathcal{I}}$ and $(s, r) \in \mathcal{T}_{\mathcal{I}\times\mathcal{I}}$. We denote the number of operations required to perform the update

$$Z|_{\hat{t}\times\hat{r}} \leftarrow Z|_{\hat{t}\times\hat{r}} + X|_{\hat{t}\times\hat{s}} Y|_{\hat{s}\times\hat{r}}$$

by $W_{\mathrm{mm}}(t, s, r)$.

Case 1: (t,s) and (s,r) are subdivided. If (t,s) and (s,r) are not leaves of $\mathcal{T}_{\mathcal{I}\times\mathcal{I}}$, we use recursion to perform updates

$$Z|_{\hat{t}'\times\hat{r}'} \leftarrow Z|_{\hat{t}'\times\hat{r}'} + X|_{\hat{t}'\times\hat{s}'} Y|_{\hat{s}'\times\hat{r}'}$$

for all $t' \in \operatorname{sons}^+(t)$, $s' \in \operatorname{sons}^+(s)$ and $r' \in \operatorname{sons}^+(r)$, where the abbreviation

$$\operatorname{sons}^+(t) := \begin{cases} \operatorname{sons}(t) & \text{if } \operatorname{sons}(t) \neq \emptyset, \\ \{t\} & \text{otherwise} \end{cases} \qquad \text{for all } t \in \mathcal{T}_{\mathcal{I}}$$

is convenient to express $\operatorname{sons}(t,s) = \operatorname{sons}^+(t) \times \operatorname{sons}^+(s)$. If (t,r) is also not a leaf of $\mathcal{T}_{\mathcal{I}\times\mathcal{I}}$, no additional algebraic operations are required and we obtain

$$W_{\mathrm{mm}}(t,s,r) = \sum_{\substack{t'\in\operatorname{sons}^+(t),\\ s'\in\operatorname{sons}^+(s),\\ r'\in\operatorname{sons}^+(r)}} W_{\mathrm{mm}}(t',s',r'). \tag{8}$$

If (t,r) is a leaf, we create temporary submatrices for each block (t',r') with $t' \in \operatorname{sons}^+(t)$ and $r' \in \operatorname{sons}^+(r)$, fill them as before using recursion, and then extend them by zero and add them to $Z|_{\hat{t}\times\hat{r}}$.

According to Theorem 5, we can find a constant C_{up} such that the low-rank update of the block $\hat{t}\times\hat{r}$ requires not more than

$$C_{\mathrm{up}} k^2 (\#\hat{t} + \#\hat{r})$$

operations and conclude

$$W_{\mathrm{mm}}(t,s,r) \leq 4C_{\mathrm{up}} k^2 (\#\hat{t} + \#\hat{r}) + \sum_{\substack{t'\in\operatorname{sons}^+(t),\\ s'\in\operatorname{sons}^+(s),\\ r'\in\operatorname{sons}^+(r)}} W_{\mathrm{mm}}(t',s',r'),$$

where the additional term is due to the fact that the four temporary submatrices for (t_1,s_1), (t_1,s_2), (t_2,s_1) and (t_2,s_2) have to be added.

Case 2: (t,s) or (s,r) is a leaf. If (s,r) is a leaf, we have $Y|_{\hat{s}\times\hat{r}} = V_s S_{sr} W_r^*$ and obtain

$$X|_{\hat{t}\times\hat{s}} Y|_{\hat{s}\times\hat{r}} = X|_{\hat{t}\times\hat{s}} V_s S_{sr} W_r^* = \underbrace{X|_{\hat{t}\times\hat{s}} V_s}_{=:\widetilde{X}_{ts}} S_{sr} W_r^*.$$

Computing $\widetilde{X}_{ts} = X|_{\hat{t}\times\hat{s}} V_s$ takes k matrix-vector multiplications with the $\mathcal{H}^2$-matrix $X|_{\hat{t}\times\hat{s}}$, and each of these requires only $\mathcal{O}(k(\#\hat{t} + \#\hat{s}))$ operations [7, Theorem 3.42]. Multiplying $\widetilde{X}_{ts}$ by S_{sr} requires $\mathcal{O}(k^2 \#\hat{t})$ operations, and the low-rank update

$$Z|_{\hat{t}\times\hat{r}} \leftarrow Z|_{\hat{t}\times\hat{r}} + (\widetilde{X}_{ts} S_{sr}) W_r^*$$

can be accomplished in not more than $C_{\mathrm{up}}k^2(\#\hat{t} + \#\hat{r})$ operations due to Theorem 5. We conclude that there is a constant C_{lb} such that

$$W_{\mathrm{mm}}(t,s,r) \leq C_{\mathrm{lb}}k^2(\#\hat{t} + \#\hat{s} + \#\hat{r})$$

holds in this case. By similar arguments we get the same estimate for the case that (t,s) is a leaf.

Matrix multiplication. Combining Case 1 and Case 2 yields a constant C_{mb} such that

$$W_{\mathrm{mm}}(t,s,r) \leq C_{\mathrm{mb}}k^2(\#\hat{t} + \#\hat{s} + \#\hat{r})$$

if $(t,s) \in \mathcal{L}_{\mathcal{I}\times\mathcal{I}}$ or $(s,r) \in \mathcal{L}_{\mathcal{I}\times\mathcal{I}}$ and

$$W_{\mathrm{mm}}(t,s,r) \leq C_{\mathrm{mb}}k^2(\#\hat{t} + \#\hat{s} + \#\hat{r}) + \sum_{\substack{t'\in\mathrm{sons}^+(t),\\ s'\in\mathrm{sons}^+(s),\\ r'\in\mathrm{sons}^+(r)}} W_{\mathrm{mm}}(t',s',r')$$

otherwise. To obtain an estimate, we follow the approach outlined in [7, Section 7.7]: we define the *triple tree* $\mathcal{T}_{\mathcal{I}\times\mathcal{I}\times\mathcal{I}}$ recursively by $\mathrm{root}(\mathcal{T}_{\mathcal{I}\times\mathcal{I}\times\mathcal{I}}) = (\mathrm{root}(\mathcal{T}_{\mathcal{I}}), \mathrm{root}(\mathcal{T}_{\mathcal{I}}), \mathrm{root}(\mathcal{T}_{\mathcal{I}}))$ and

$$\mathrm{sons}(t,s,r) := \begin{cases} \emptyset & \text{if } (t,s) \in \mathcal{L}_{\mathcal{I}\times\mathcal{I}} \text{ or } (s,r) \in \mathcal{L}_{\mathcal{I}\times\mathcal{I}}, \\ \mathrm{sons}^+(t)\times\mathrm{sons}^+(s)\times\mathrm{sons}^+(r) & \text{otherwise} \end{cases}$$
$$\text{for all } (t,s,r) \in \mathcal{T}_{\mathcal{I}\times\mathcal{I}\times\mathcal{I}}$$

and find the estimate

$$W_{\mathrm{mm}}(t_0,s_0,r_0) \leq C_{\mathrm{mb}}k^2 \sum_{(t,s,r)\in\mathcal{T}_{\hat{t}_0\times\hat{s}_0\times\hat{r}_0}} (\#\hat{t} + \#\hat{s} + \#\hat{r}) \quad \text{for all } (t_0,s_0,r_0) \in \mathcal{T}_{\mathcal{I}\times\mathcal{I}\times\mathcal{I}},$$

where $\mathcal{T}_{\hat{t}_0\times\hat{s}_0\times\hat{r}_0}$ denotes the subtree of $\mathcal{T}_{\mathcal{I}\times\mathcal{I}\times\mathcal{I}}$ with the root (t_0,s_0,r_0).

As in [7, Lemma 8.8], $(t,s,r) \in \mathcal{T}_{\mathcal{I}\times\mathcal{I}\times\mathcal{I}}$ implies $(t,s) \in \mathcal{T}_{\mathcal{I}\times\mathcal{I}}$ and $(s,r) \in \mathcal{T}_{\mathcal{I}\times\mathcal{I}}$, so we can prove that each t, s or $r \in \mathcal{T}_{\mathcal{I}}$ appears in not more than C_{sp}^2 triples in $\mathcal{T}_{\mathcal{I}\times\mathcal{I}\times\mathcal{I}}$. Denoting the depth of the cluster tree $\mathcal{T}_{\mathcal{I}}$ by $p_{\mathcal{I}}$ and using the standard estimate

$$\sum_{t\in\mathcal{T}_{\hat{t}_0}} \#\hat{t} \leq (p_{\mathcal{I}}+1)\#\hat{t}_0$$

(cf., e.g., [7, Corollary 3.10]), we obtain the following result:

Theorem 6 (Complexity, multiplication) *Let $p_{\mathcal{I}}$ denote the depth of the cluster tree $\mathcal{T}_{\mathcal{I}}$. We have*

$$W_{\mathrm{mm}}(t_0,s_0,r_0) \leq C_{\mathrm{mm}}k^2(p_{\mathcal{I}}+1)(\#\hat{t}_0 + \#\hat{s}_0 + \#\hat{r}_0) \quad \textit{for all } (t_0,s_0,r_0) \in \mathcal{T}_{\mathcal{I}\times\mathcal{I}\times\mathcal{I}}$$

with $C_{\mathrm{mm}} := C_{\mathrm{sp}}^2 C_{\mathrm{mb}}$.

Forward substitution Let us now consider the forward substitution, e.g., solving

$$L|_{\hat{t}\times\hat{t}}X|_{\hat{t}\times\hat{s}} = Y|_{\hat{t}\times\hat{s}}$$

for a block $(t,s) \in \mathcal{T}_{\mathcal{I}\times\mathcal{I}}$. We denote the number of operations required by $W_{\mathrm{lfs}}(t,s)$.

If t is a leaf, we perform a forward substitution for a standard matrix, which takes half the operations of the matrix multiplication, i.e., we have

$$W_{\mathrm{lfs}}(t,s) \leq W_{\mathrm{mm}}(t,t,s)/2 \leq W_{\mathrm{mm}}(t,t,s). \tag{9a}$$

If t is not a leaf, we have $\operatorname{sons}(t) = \{t_1, t_2\}$ and our algorithm performs the following steps

- solve $L|_{\hat{t}_1\times\hat{s}'}X|_{\hat{t}_1\times\hat{s}'} = Y|_{\hat{t}_1\times\hat{s}'}$ for all $s' \in \operatorname{sons}^+(s)$,
- compute $\widetilde{Y}|_{\hat{t}_2\times\hat{s}'} := Y|_{\hat{t}_2\times\hat{s}} - L|_{\hat{t}_2\times\hat{t}_1}X|_{\hat{t}_1\times\hat{s}'}$ for all $s' \in \operatorname{sons}^+(s)$, and
- solve $L|_{\hat{t}_2\times\hat{t}_2}X|_{\hat{t}_2\times\hat{s}'} = \widetilde{Y}|_{\hat{t}_2\times\hat{s}'}$ for all $s' \in \operatorname{sons}^+(s)$.

In total, we require

$$W_{\mathrm{lfs}}(t,s) \leq \sum_{s'\in\operatorname{sons}^+(s)} W_{\mathrm{lfs}}(t_1,s') + W_{\mathrm{mm}}(t_2,t_1,s') + W_{\mathrm{lfs}}(t_2,s') \tag{9b}$$

operations. Solving this recurrence relation yields the following estimate:

Theorem 7 (Complexity, forward substitution) *We have*

$$W_{\mathrm{lfs}}(t,s) \leq W_{\mathrm{mm}}(t,t,s) \qquad \text{for all } (t,s) \in \mathcal{T}_{\mathcal{I}\times\mathcal{I}}. \tag{10}$$

Proof. By structural induction in $t \in \mathcal{T}_{\mathcal{I}}$.

If $t \in \mathcal{T}_{\mathcal{I}}$ is a leaf, (10) follows directly from (9a).

If $t \in \mathcal{T}_{\mathcal{I}}$ is not a leaf, let $\operatorname{sons}(t) = \{t_1, t_2\}$ and assume that (10) holds for (t_1, s') and (t_2, s') for all $s' \in \operatorname{sons}^+(s)$. With (9b), we get

$$\begin{aligned}
W_{\mathrm{lfs}}(t,s) &\leq \sum_{s'\in\operatorname{sons}^+(t)} W_{\mathrm{lfs}}(t_1,s') + W_{\mathrm{mm}}(t_2,t_1,s') + W_{\mathrm{lfs}}(t_2,s') \\
&\leq \sum_{s'\in\operatorname{sons}^+(t)} W_{\mathrm{mm}}(t_1,t_1,s') + W_{\mathrm{mm}}(t_2,t_1,s') + W_{\mathrm{mm}}(t_2,t_2,s') \\
&\leq \sum_{s'\in\operatorname{sons}^+(t)} W_{\mathrm{mm}}(t_1,t_1,s') + W_{\mathrm{mm}}(t_2,t_1,s') + W_{\mathrm{mm}}(t_1,t_2,s') + W_{\mathrm{mm}}(t_2,t_2,s') \\
&= \sum_{t'\in\operatorname{sons}^+(t)} \sum_{t''\in\operatorname{sons}^+(t)} \sum_{s'\in\operatorname{sons}^+(s)} W_{\mathrm{mm}}(t',t'',s') \leq W_{\mathrm{mm}}(t,t,s).
\end{aligned}$$

This completes the induction. □

Our algorithm for solving $X|_{\hat{t}\times\hat{s}}R|_{\hat{s}\times\hat{s}} = Y|_{\hat{t}\times\hat{s}}$ can be treated by the same arguments to find that the number of operations $W_{\mathrm{rfs}}(t,s)$ is bounded by $W_{\mathrm{mm}}(t,s,s)$.

LR decomposition Finally we consider the LR factorization, e.g., finding $L|_{\hat{t}\times\hat{t}}, R|_{\hat{t}\times\hat{t}}$ such that

$$L|_{\hat{t}\times\hat{t}} R|_{\hat{t}\times\hat{t}} = A|_{\hat{t}\times\hat{t}}.$$

We denote the complexity by $W_{\mathrm{lr}}(t)$.

If t is a leaf, computing the LR factorization directly requires one third of the number of operations required to compute the matrix product, i.e., we have

$$W_{\mathrm{lr}}(t) \leq W_{\mathrm{mm}}(t,t,t)/3 \leq W_{\mathrm{mm}}(t,t,t). \tag{11a}$$

If t is not a leaf, i.e., if $\operatorname{sons}(t) = \{t_1, t_2\}$ holds, our recursive algorithm performs the following steps:

- compute the factorization of $A|_{\hat{t}_1\times\hat{t}_1}$,
- solve $L|_{\hat{t}_1\times\hat{t}_1} R|_{\hat{t}_1\times\hat{t}_2} = A|_{\hat{t}_1\times\hat{t}_2}$ for $R|_{\hat{t}_1\times\hat{t}_2}$,
- solve $L|_{\hat{t}_2\times\hat{t}_1} R|_{\hat{t}_1\times\hat{t}_1} = A|_{\hat{t}_2\times\hat{t}_1}$ for $L|_{\hat{t}_2\times\hat{t}_1}$,
- compute $\widetilde{A}|_{\hat{t}_2\times\hat{t}_2} = A|_{\hat{t}_2\times\hat{t}_2} - L|_{\hat{t}_2\times\hat{t}_1} R|_{\hat{t}_1\times\hat{t}_2}$,
- compute the factorization of $\widetilde{A}|_{\hat{t}_2\times\hat{t}_2}$,

and this takes

$$\begin{aligned} W_{\mathrm{lr}}(t) &= W_{\mathrm{lr}}(t_1) + W_{\mathrm{lfs}}(t_1,t_2) + W_{\mathrm{rfs}}(t_2,t_1) + W_{\mathrm{mm}}(t_2,t_1,t_2) + W_{\mathrm{lr}}(t_2) \\ &\leq W_{\mathrm{lr}}(t_1) + W_{\mathrm{mm}}(t_1,t_1,t_2) + W_{\mathrm{mm}}(t_2,t_1,t_1) + W_{\mathrm{mm}}(t_2,t_1,t_2) + W_{\mathrm{lr}}(t_2) \end{aligned} \tag{11b}$$

operations due to Theorem 7. We can solve this recurrence relation to obtain the following result:

Theorem 8 (Complexity, LR factorization) *We have*

$$W_{\mathrm{lr}}(t) \leq W_{\mathrm{mm}}(t,t,t) \qquad \text{for all } t \in \mathcal{T}_{\mathcal{I}}. \tag{12}$$

Proof. By structural induction in $t \in \mathcal{T}_{\mathcal{I}}$.

If $t \in \mathcal{T}_{\mathcal{I}}$ is a leaf, (12) follows directly from (11a).

If $t \in \mathcal{T}_{\mathcal{I}}$ is not a leaf, we have $\operatorname{sons}(t) = \{t_1, t_2\}$ and assume that (12) holds for t_1 and t_2. With (11b), we get

$$\begin{aligned} W_{\mathrm{lr}}(t) &\leq W_{\mathrm{lr}}(t_1) + W_{\mathrm{mm}}(t_1,t_1,t_2) + W_{\mathrm{mm}}(t_2,t_1,t_1) + W_{\mathrm{mm}}(t_2,t_1,t_2) + W_{\mathrm{lr}}(t_2) \\ &\leq W_{\mathrm{mm}}(t_1,t_1,t_1) + W_{\mathrm{mm}}(t_1,t_1,t_2) + W_{\mathrm{mm}}(t_2,t_1,t_1) \\ &\quad + W_{\mathrm{mm}}(t_2,t_1,t_2) + W_{\mathrm{mm}}(t_2,t_2,t_2) \\ &\leq \sum_{\substack{t'\in\operatorname{sons}^+(t),\\ t''\in\operatorname{sons}^+(t),\\ t'''\in\operatorname{sons}^+(t)}} W_{\mathrm{mm}}(t',t'',t''') \leq W_{\mathrm{mm}}(t,t,t). \end{aligned}$$

This completes the induction. □

Now we can state the final result for the complexity of the approximation of the LR factorization:

Corollary 9 (Global complexity) *The approximate LR factorization of rank k can be constructed in $\mathcal{O}(k^2(p_{\mathcal{I}}+1)n_{\mathcal{I}})$ operations, where $p_{\mathcal{I}}$ denotes the depth of the cluster tree.*

Proof. Combine Theorem 6 with Theorem 8 and apply it to $t = \operatorname{root}(\mathcal{T}_{\mathcal{I}})$. □

Remark 10 (Inversion) *By a similar approach, we can also obtain an approximation of the inverse: consider $A|_{\hat{t}\times\hat{t}}^{-1}$ for $t \in \mathcal{T}_{\mathcal{I}}$. If t is a leaf, we can compute the inverse directly. Otherwise we let*

$$A|_{\hat{t}\times\hat{t}} = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}$$

as before and perform a block Gauss elimination to obtain

$$A|_{\hat{t}\times\hat{t}} = \begin{pmatrix} I & \\ A_{21}A_{11}^{-1} & I \end{pmatrix} \begin{pmatrix} A_{11} & A_{12} \\ & A_{22} - A_{21}A_{11}^{-1}A_{12} \end{pmatrix}.$$

We denote the Schur complement by $S := A_{22} - A_{21}A_{11}^{-1}A_{12}$, compute the inverses of the two block triangular matrices, and find the representation

$$A|_{\hat{t}\times\hat{t}}^{-1} = \begin{pmatrix} A_{11}^{-1} & -A_{11}^{-1}A_{12}S^{-1} \\ & S^{-1} \end{pmatrix} \begin{pmatrix} I & \\ -A_{21}A_{11}^{-1} & I \end{pmatrix}$$

of the inverse. This equation allows us to compute the inverse $A|_{\hat{t}\times\hat{t}}^{-1}$ by recursively computing A_{11}^{-1} and S^{-1} and carrying out six matrix multiplications. Following the same reasoning as before, we can prove that the number of operations $W_{\mathrm{inv}}(t)$ is bounded by $W_{\mathrm{mm}}(t,t,t)$, and Theorem 6 yields that $\mathcal{O}(k^2(p_{\mathcal{I}}+1)n_{\mathcal{I}})$ operations are sufficient to construct the approximate inverse.

7 Numerical experiments

We investigate the practical properties of the new algorithms by considering two standard model problems: for the first model problem, we consider the linear system resulting from a finite element discretization of Poisson's equation on the unit square using piecewise linear nodal basis functions on a regular mesh. We use the domain decomposition cluster strategy similar to the one described in [14] to find a suitable cluster tree and block tree for the sparse system. Since the matrix is symmetric and positive definite, we construct an approximate Cholesky factorization by the algorithms described in Section 3 using a block-relative accuracy of $\hat{\epsilon} \in \mathbb{R}_{>0}$ for the recompression.

Table 1 lists the results for grid levels 7 to 12 with 16 129 to 16 769 025 grid points. The column "Param" gives the admissibility parameter η and the relative accuracy $\hat{\epsilon}$ used in the recompression algorithm, the column "Setup" gives the time (in seconds per degree of freedom) and storage requirements (in KB per degree of freedom) for constructing the Cholesky factorization, and the column "Solve" gives an approximation of the convergence factor $\|I - \widetilde{A}^{-1}A\|_2$ obtained by the power iteration, the number m of

Grid		Param		Setup		Solve		
ℓ	n	η	$\hat{\epsilon}$	Time/n	Mem/n	$\|I - \widetilde{A}^{-1}A\|_2$	m	Time/n
7	16 129	4	3.1_{-3}	7.0_{-5}	1.0	6.2_{-2}	3	1.9_{-6}
8	65 025	4	7.7_{-4}	8.4_{-5}	1.1	6.9_{-2}	3	2.3_{-6}
9	261 121	4	1.9_{-4}	1.1_{-4}	1.2	6.8_{-2}	3	2.6_{-6}
10	1 046 529	4	4.8_{-5}	1.3_{-4}	1.2	1.0_{-1}	3	2.8_{-6}
11	4 190 209	4	1.2_{-5}	1.5_{-4}	1.2	1.1_{-1}	3	3.0_{-6}
12	16 769 025	4	3.0_{-6}	1.7_{-4}	1.2	9.6_{-2}	3	3.5_{-6}

Table 1: Preconditioner for the finite element model problem.

Grid		Param		Setup		Solve		
ℓ	n	η	$\hat{\epsilon}$	Time/n	Mem/n	$\|I - \widetilde{A}^{-1}A\|_2$	m	Time/n
11	8 192	2	1.2_{-4}	1.6_{-4}	0.8	4.5_{-2}	3	1.6_{-6}
12	16 384	2	6.1_{-5}	1.9_{-4}	0.8	2.7_{-2}	3	1.6_{-6}
13	32 768	2	3.1_{-5}	2.1_{-4}	0.8	2.9_{-2}	3	1.7_{-6}
14	65 536	2	1.5_{-5}	2.4_{-4}	0.8	1.9_{-2}	3	1.8_{-6}
15	131 072	2	7.6_{-6}	2.6_{-4}	0.8	1.6_{-2}	3	1.9_{-6}
16	262 144	2	3.8_{-6}	3.0_{-4}	0.8	1.9_{-2}	3	2.0_{-6}
17	524 288	2	1.9_{-6}	3.3_{-4}	0.8	1.9_{-2}	3	2.1_{-6}
18	1 048 576	2	9.5_{-7}	3.5_{-4}	0.8	1.3_{-2}	4	2.2_{-6}
19	2 097 152	2	4.8_{-7}	3.6_{-4}	0.8	2.1_{-2}	3	2.2_{-6}
20	4 194 304	2	2.4_{-7}	3.8_{-4}	0.8	2.0_{-1}	4	2.2_{-6}

Table 2: Preconditioner for the boundary element model problem.

iteration steps required by the conjugate gradient iteration to reduce the relative residual norm below 10^{-8}, and the time required per step and degree of freedom.

We can see that choosing the accuracy $\hat{\epsilon}$ like $\mathcal{O}(h^2)$ to keep up with the $\mathcal{O}(h^{-2})$ growth of the condition number is sufficient to obtain a stable convergence rate. The time per degree of freedom grows like $\log(n)$, the storage requirements are bounded.

For the second model problem, we consider the linear system resulting from the finite element discretization of the single layer potential operator on the two-dimensional unit circle using a regular mesh and piecewise constant basis functions.

Table 2 lists the results for this model problem. Choosing $\hat{\epsilon}$ like $\mathcal{O}(h)$ compensates for the $\mathcal{O}(h^{-1})$ growth of the condition number and leads to stable convergence rates. As in the case of the partial differential equation, the time per degree of freedom grows like $\log(n)$ and the storage requirements per degree of freedom are bounded.

We can conclude that the new preconditioner requires a setup time of $\mathcal{O}(n \log n)$ and $\mathcal{O}(n)$ units of storage to ensure stable h-independent convergence of the conjugate gradient method.

References

[1] M. Bebendorf. Hierarchical LU decomposition based preconditioners for BEM. *Computing*, 74:225–247, 2005.

[2] M. Bebendorf. Why finite element discretizations can be factored by triangular hierarchical matrices. *SIAM J. of Numer. Anal.*, 45(4):1472–1494, 2007.

[3] M. Bebendorf and W. Hackbusch. Existence of $\mathcal{H}$-matrix approximants to the inverse FE-matrix of elliptic operators with L^∞-coefficients. *Numer. Math.*, 95:1–28, 2003.

[4] S. Börm. $\mathcal{H}^2$-matrix arithmetics in linear complexity. *Computing*, 77(1):1–28, 2006.

[5] S. Börm. Data-sparse approximation of non-local operators by $\mathcal{H}^2$-matrices. *Lin. Alg. Appl.*, 422:380–403, 2007.

[6] S. Börm. Approximation of solution operators of elliptic partial differential equations by $\mathcal{H}$- and $\mathcal{H}^2$-matrices. *Numer. Math.*, 115(2):165–193, 2010.

[7] S. Börm. *Efficient Numerical Methods for Non-local Operators: $\mathcal{H}^2$-Matrix Compression, Algorithms and Analysis*, volume 14 of *EMS Tracts in Mathematics*. EMS, 2010.

[8] S. Börm and W. Hackbusch. Data-sparse approximation by adaptive $\mathcal{H}^2$-matrices. *Computing*, 69:1–35, 2002.

[9] S. Chandrasekaran, M. Gu, and W. Lyons. A fast adaptive solver for hierarchically semiseparable representations. *Calcolo*, 42:171–185, 2005.

[10] M. Faustmann, J. M. Melenk, and D. Praetorius. Existence of $\mathcal{H}$-matrix approximants to the inverses of BEM matrices: the simple-layer operator. Technical Report 37, Institut für Analysis und Scientific Computing, TU Wien, 2013. available at `http://www.asc.tuwien.ac.at/preprint/2013/asc37x2013.pdf`.

[11] M. Faustmann, J. M. Melenk, and D. Praetorius. $\mathcal{H}$-matrix approximability of the inverse of FEM matrices. Technical Report 20, Institut für Analysis und Scientific Computing, TU Wien, 2013. available at `http://www.asc.tuwien.ac.at/preprint/2013/asc20x2013.pdf`.

[12] L. Grasedyck. Adaptive recompression of $\mathcal{H}$-matrices for BEM. *Computing*, 74(3):205–223, 2004.

[13] L. Grasedyck and W. Hackbusch. Construction and arithmetics of $\mathcal{H}$-matrices. *Computing*, 70:295–334, 2003.

[14] L. Grasedyck, R. Kriemann, and S. LeBorne. Domain decomposition based $\mathcal{H}$-LU preconditioning. *Numer. Math.*, 112(4):565–600, 2009.

[15] W. Hackbusch. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part I: Introduction to $\mathcal{H}$-matrices. *Computing*, 62:89–108, 1999.

[16] W. Hackbusch. *Hierarchische Matrizen — Algorithmen und Analysis*. Springer, 2009.

[17] W. Hackbusch and B. N. Khoromskij. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part II: Application to multi-dimensional problems. *Computing*, 64:21–47, 2000.

[18] W. Hackbusch, B. N. Khoromskij, and S. A. Sauter. On $\mathcal{H}^2$-matrices. In H. Bungartz, R. Hoppe, and C. Zenger, editors, *Lectures on Applied Mathematics*, pages 9–29. Springer-Verlag, Berlin, 2000.

[19] M. Lintner. The eigenvalue problem for the 2d Laplacian in $\mathcal{H}$-matrix arithmetic and application to the heat and wave equation. *Computing*, 72:293–323, 2004.

[20] P. G. Martinsson. A fast direct solver for a class of elliptic partial differential equations. *J. Sci. Comput.*, 38:316–330, 2008.

[21] P. G. Martinsson, V. Rokhlin, and M. Tygert. A fast algorithm for the inversion of general Toeplitz matrices. *Comp. Math. Appl.*, 50:741–752, 2005.

[22] V. Rokhlin. Rapid solution of integral equations of classical potential theory. *J. Comp. Phys.*, 60:187–207, 1985.

[23] J. Xia, S. Chandrasekaran, M. Gu, and X. Li. Superfast multifrontal method for large structured linear systems of equations. *SIAM J. Matrix Anal. Appl.*, 31(3):1382–1411, 2009.

[24] J. Xia, S. Chandrasekaran, M. Gu, and X. S. Li. Fast algorithms for hierarchically semiseparable matrices. *Numer. Lin. Alg. Appl.*, 2009. available at `http://dx.doi.org/10.1002/nla.691`.